

DTIC FILE COPY

(4)

REPORT DOCUMENTATION PAGE

AD-A200 695

2b. DECLASSIFICATION/DOWNGRADING SCHEDULE		1d. RESTRICTIVE MARKINGS	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) LIDS-FR-1821		3. DISTRIBUTION/AVAILABILITY OF REPORT Unlimited	
6a. NAME OF PERFORMING ORGANIZATION Massachusetts Institute of Technology	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Office of Naval Research	
6c. ADDRESS (City, State, and ZIP Code) 77 Massachusetts Avenue Cambridge, Massachusetts 02139		7b. ADDRESS (City, State, and ZIP Code) 800 North Quincy Street Arlington, Virginia 22217	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Office of Naval Research	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER	
8c. ADDRESS (City, State, and ZIP Code) 800 North Quincy Street Arlington, Virginia 22217		10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO. PROJECT NO. 95006 TASK NO. WORK UNIT ACCESSION NO.	
11. TITLE (Include Security Classification) Data Networks Reliability			
12. PERSONAL AUTHOR(S) Robert G. Gallager			
13a. TYPE OF REPORT FINAL	13b. TIME COVERED FROM 1984 TO 1986	14. DATE OF REPORT (Year, Month, Day) October 3, 1988	15. PAGE COUNT 16
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES FIELD GROUP SUB-GROUP		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) The research from 1984 to 1986 on Data Network Reliability had the objective of developing general principles governing the reliable and efficient control of data networks. The research was centered around three major areas: congestion control, multiaccess networks, and distributed asynchronous algorithms. The major topics within congestion control were the use of flow control to reduce congestion and the use of routing to reduce congestion. The major topics within multiaccess networks were the communication properties of multiaccess channels, collision resolution, and packet radio networks. The major topics within asynchronous distributed algorithms were failure recovery, time vs. communication tradeoffs, and the general theory of distributed algorithms.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Robert D. Nelson		22b. TELEPHONE (Include Area Code) 22c. OFFICE SYMBOL (617) 253-2533	

FINAL REPORT

DATA NETWORK RELIABILITY

ABSTRACT

The research from 1984 to 1986 on Data Network Reliability had the objective of developing general principles governing the reliable and efficient control of data networks. The research was centered around three major areas: congestion control, multiaccess networks, and distributed asynchronous algorithms. The major topics within congestion control were the use of flow control to reduce congestion and the use of routing to reduce congestion. The major topics within multiaccess networks were the communication properties of multiaccess channels, collision resolution, and packet radio networks. The major topics within asynchronous distributed algorithms were failure recovery, time vs. communication tradeoffs, and the general theory of distributed algorithms.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



88 10 18 130

FINAL REPORT

DATA NETWORK RELIABILITY

TABLE OF CONTENTS

	Page
Abstract	1
1. Objectives and Summary	2
1.1 Brief Statement of Objectives	2
1.2 Summary of Results	2
2. Detailed Results	4
2.1 Results on Flow Control and Routing	4
2.1.2 Routing	7
2.2 Multiaccess Networks	8
2.2.1 Multiaccess Communication Channels	8
2.2.2 Random Access and Collision Resolution	10
2.2.3 Packet Radio Multiaccess Systems	10
2.3 Distributed Algorithms	11
2.3.1 Failure Recovery Algorithms	11
2.3.2 Other Distributed Algorithms And Protocols	12
References	

1. OBJECTIVES AND SUMMARY

1.1 Brief Statement of Objectives

Military operations will depend increasingly on data communication networks as weapon systems and command and control systems become more automated and require faster response times. The objectives of this research were to discover and develop a quantitative understanding of the fundamental principles governing the reliability and control of the data networks themselves. Particular emphasis was given to responsive network control to compensate for rapidly changing traffic requirements and node and link performance capabilities.

1.2 Summary of Results

Our research was centered around three major areas: congestion control by means of flow control and routing; networks with multiple access channels, including packet radio, local area networks, coding, and spread spectrum; and finally asynchronous distributed algorithms, including failure recovery algorithms, time vs. communication tradeoffs, and algorithms for broadcast message dissemination.

In the area of flow control, the emphasis was on achieving a balance between fairness (subject to priorities) and small delay. The main results here are documented in [Mos84], [Hah86], [HaG86], [Muk86], [Reg86], [ReH87a,b], [Tha84], [BeG87]. In the area of routing, the emphasis was on developing asynchronous algorithms with various types of optimality properties; this is documented in [BeB84] and [TsB86].

In the area of multiple access networks, one major topic centered on the communication properties of the multiaccess channels; this work is documented in

[Gal85a], [Hui84a,b], [HuH85], [Hum85b], [Ari85]. Another major topic was Random access and collision resolution, as documented in [Hum85a], [Yat86], [Kha85], and [BeG87]. The last topic was packet radio; the results here are documented in [Par85], [Ros86], [Mar84], [Esc85], [BeG87].

In the area of asynchronous distributed algorithms, the topic of failure recovery has been of fundamental interest to us for many years; our results over the contract period are documented in [Spi85], [Spi86], [SoH87], [Cas86], and [BeG87]. Our research on the general methodology of distributed algorithms is documented in [AwG85a,b], Hum84], and [Gal85b].

2. DETAILED RESULTS

A major accomplishment during the period of the grant was the completion of a text book on Data Networks [BeG87]. This text provides a major synthesis of our work and that of others, and should help to provide a more unified focus for future research and development in the data network field. Although the grant did not support the actual preparation of the book, the grant did support much of the basic research that allowed the structure of the network field to be presented within a cohesive framework. More specific research accomplishments are as follows:

2.1 Results on Flow Control and Routing

2.1.1 Flow control: Flow control is the set of mechanisms used in a network to maintain equitable and acceptable delay for all users. Part of the function of flow control is to decide when to accept and when to reject packets into the subnetwork. Such a decision involves both fairness between the users and the tradeoff between the undesirability of rejecting packets and the undesirability of subjecting them to long delays within the subnetwork. A second function of flow control is to choose the order in which packets are transmitted on the links of the subnetwork; this order is often first-come-first-serve (FCFS) by default, but other possibilities can reduce delay or increase fairness between the users. A third function is to prevent deadlock when the buffers at a node fill up. Most of our work has been done under the assumption of infinite buffer space, where the third problem does not arise. Our premise is that by concentrating on small delay, the buffer utilization will also be small. This means that with suitably large buffers, buffer overflow will be a rare situation. Thus the problem of preventing deadlock when buffers become full

becomes a distributed algorithm question that is separable from the questions of delay and fairness.

In the early part of the contract, much of our flow control work focused on max-min flow control. The criterion of fairness here is that the data rate R_i for each user i should be such that any increase in R_i for user i must lead to a decrease for some user j for which $R_j \leq R_i$. This criterion leads to a unique assignment of rates for all users in a network with given routes. The max-min criterion can also be easily generalized to fairness between different classes of users with different requirements. One of the major problems with this criterion is that of translating the notion of rate into dynamic questions of network control, such as admitting packets into the network and ordering their service. For integrated networks serving voice and data, however, the rate of a voice user is a meaningful quantity which can be controlled in the voice digitization process. Even for voice applications, however, a distributed algorithm is necessary to continually update max-min rates as the set of users in the network changes.

J. Mosely [Mos84], in a Ph. D. thesis, developed distributed algorithms for calculating max-min fair rates. These algorithms have the property that they converge to correct values for all the rates, and this convergence takes place in the presence of arbitrary delays in the communication used within the distributed algorithm. Previous algorithms could be shown to converge only when all the updates are synchronous or when the time between successive updates is substantially longer than the propagation delays across the network.

Another approach to achieving the max-min fairness criterion is to use a round robin service discipline at each link, sending a packet for the first session on a link, then one for the second session, and so forth to the last session before returning to the first session again. If there are no packets waiting from a given session, the algorithm immediately goes on to the next session. This round robin strategy ensures that if a session

always has at least one waiting packet for a given link, then that session receives as much service as any other session using that link.

E. Hahne [Hah86], in a Ph. D. thesis, shows that if a round robin service discipline is combined with link by link window flow control of sufficiently large window size, then the resulting rates are fair according to the max-min criterion. This result has two important implications. The first is that max-min fairness can be easily achieved without complex algorithms to calculate max-min rates. The second is that in situations where session rates are not meaningful (such as for bursty sources), one still maintains the essential idea of fairness - namely, whenever a session has data waiting for service at a link, that session gets as much service as any other user on that link. The window sizes required to absolutely guarantee max-min fairness are impractically large. The thesis also demonstrates, however, that if a session uses small windows, it is guaranteed a small delay, perhaps at some degradation in maximum throughput.

Another approach to the use of service disciplines at the nodes to achieve small delay and fairness between users was investigated in an Sc. D. thesis by U. Mukherji [Muk86]. In this approach, the transmission on each link is broken into frames, where a frame consists of one packet slot for each session using the link. If there are no packets waiting for the given session, that slot is used by some other session on a FCFS basis subject to an end to end windowing strategy. This strategy provides each session with a guaranteed throughput and a guaranteed maximum delay (where the maximum delay increases with window size). Such guarantees are desirable for packet voice, whereas the possibility of extra throughput when the network is lightly loaded is desirable for data traffic.

J. Regnier [Reg86], [ReH87a,b], in his Ph. D. thesis, studied the problem of allocating queueing priorities to virtual circuits sharing a link, in order to minimize some system functions. This is distinct from the usual approach of assigning priorities and then

evaluating performances. He was able to completely characterize feasible multidimensional throughput/delay regions in which the underlying queuing discipline is non-preemptive and inputs are Poisson. Although the feasible region has exponentially many extremal points, he was able to exploit its structure and to develop a fast algorithm to assign priorities. His work also extends to max-min criteria similar to those outlined above. With some independence assumptions, his work was extended to networks. He developed distributed asynchronous algorithms to allow the links in a network to cooperate in assigning queueing priorities to the virtual circuits established over them. In some more restrictive cases he also characterized optimal rate/priority assignments and developed algorithms to achieve them. An earlier somewhat similar attempt to reduce delay by scheduling was done by Thabit [Tha84].

2.1.2 Routing: Most of the existing analysis of distributed routing algorithms is predicated on several assumptions that are to some extent violated in practice. These are as follows.

- 1) The *quasi-static assumption*, i.e., the external traffic arrival process for each OD pair is stationary over time. This assumption is approximately valid when there is a large number of user-pair conversations associated with each OD pair, and each of these conversations has an arrival rate that is small relative to the total arrival rate for the OD pair (i.e., a "many small users" assumption). We have provided in [1] an asymptotic analysis of the effect of violation of this assumption on the stationary character of the external traffic arrival rates.

- 2) The *fast settling time assumption*, i.e., transients in the link flows due to changes in routing are negligible. In other words, once the routing is updated, the link flows settle to their new values within time which is very small relative to the time between routing updates. This assumption is typically valid in datagram networks but less so in

virtual circuit networks where existing virtual circuits may not be rerouted after a routing update. When this assumption is violated, link flow measurements reflect a dependence not just on the current routing but also possible on several past routings.

3) The *synchronous update assumption*, i.e., all link flows are measured simultaneously, and are received simultaneously at all network nodes who in turn simultaneously carry out a routing update. However, there may be technical reasons (such as software complexity) that argue against enforcing a synchronous update protocol. For example, the distributed routing algorithm of the ARPANET is not operated synchronously. Furthermore, in an asynchronous updating environment, the link flows are typically measured as time averages that reflect dependence on more than one update.

In the paper [2] we have studied routine methods that are valid even if the settling time and synchronous update assumption are violated to a considerable extent. Even though we retain the quasi-static assumption in our analysis, we conjecture that the result of this paper can be generalized along the lines of another related study [3] where it is shown that a routing algorithm based on a shortest path rule converges to a neighborhood of the optimum. The size of this neighborhood depends on the extent of the violation of the quasi-static assumption. A similar deviation from optimality can be caused by errors in the measurement of the link flows. In our analysis, these errors are neglected.

2.2 Multiaccess Networks

2.2.1 Multiaccess Communication Channels: A review paper [Gal85a] providing some perspective on multiaccess communication was written. This contrasts the random access approach, which is based on transmitting one packet at a time and resolving collisions if more than one transmission occurs, with the information theoretic and spread spectrum approach in which many packets are transmitted simultaneously.

J. Hui, [Hui84a,b], has combined the results of multiaccess information theory with those of classical communication theory to treat many aspects of coding for multiaccess channels, including asynchronism between transmitters and absence of a priori knowledge at the receiver about which transmitters are active. He was able to find the achievable rate region for a number of simple multiaccess models such as the collision model and a spread spectrum model. In addition he found effective coding strategies for some of these channels, e.g., one that achieves a throughput of 0.29 packets per slot on the collision channel without feedback. In an important theoretical development, he and Humblet [HuH85] showed that if there is no synchronization between transmitters, then the achievable region of multiaccess information theory is reduced; those rates that are achievable only as convex combinations of directly achievable rates are not achievable in the asynchronous case. Earlier, it had been shown in [CMP80] that the full achievable region is achievable if there is only a bounded degree of asynchronism.

E. Arikan, in a Ph.D. thesis [Ari85], extended sequential decoding for use on multiaccess channels. It was discovered that simple modifications of the classical Fano metric do not operate satisfactorily for multiaccess channels, so a new type of metric was developed. The computational cut-off rate for sequential decoding was also generalized to a cut-off rate region, and the complexity and effectiveness of joint decoding versus independent decoding was analyzed. Finally, it was shown that the cut-off rate is actually the maximum achievable rate for sequential decoding for a single source, thus generalizing the classic result of Jacobs and Berlekamp [JaB67] and settling a 20 year old conjecture.

Finally, the paradox (see [McE81], [Mas81]) concerning infinite capacity and finite computational cut off rate for the classical model of the point to point direct detection optical channel was fully explained by Humblet [Hum85b]. This result is generalized in the appendix to an entire class of channels that have infinite capacity with an energy constraint and an unbounded number of degrees of freedom.

2.2.2 Random Access and Collision Resolution: One of the difficulties with the splitting algorithms that achieve maximum throughput on collision channels is that they require all transmitters to monitor the feedback information from the receiver at all times. Humblet [Hum85a] developed a new collision resolution algorithm that transmits packets in a last-come-first-serve order (see also [GeP85]). In this new algorithm, transmitters need only monitor the feedback when they have packets to send, and the throughput can be arbitrarily close to that of the best splitting algorithms, namely 0.487 packets per slot.

One of the classic results on collision channels, established elsewhere by N. Pippenger [Pip81], is that if the receiver feedback identifies the number of transmitters involved in a collision, then throughputs arbitrarily close to one packet per slot can be achieved. R. Yates [Yat86], in an S.M. thesis, investigated the delay involved in such strategies and showed that delay increases so quickly with throughput as to make such strategies impractical.

The delay incurred by making reservations for the use of a satellite multiaccess channel was analyzed by A. Khanna [Kha85] in another S. M. thesis. He used part of the band for making reservations by FDM and found the best split between reservation bandwidth and data bandwidth (see also [BeG87] for a more refined analysis).

2.2.3 Packet Radio Multiaccess Systems: One approach to achieving high throughput in a packet radio network is to organize the network into a partition of clusters in which each cluster has a central node with a direct link to all other nodes of the cluster; inter-cluster communication then takes place via gateway nodes connecting the clusters (see [BaE81]). It is desirable to minimize the number of clusters so as to simplify the inter-cluster communication. A. Parekh [Par85], in an S. M. thesis, showed that the minimum

cluster problem is NP complete; he went on, however, to develop several new distributed algorithms that are quite effective in partitioning a net into a small number of clusters.

The problem of choosing power levels for transmitters in a packet radio net (or equivalently choosing transmission radii) has been studied for many years (see [TaK84] for an up to date analysis). C. Rossi [Ros86], in an S. M. thesis, has investigated this problem under the assumptions of a slotted system and perfect capture; that is, if more than one transmitter is active at a time, a receiver will successfully receive the packet from the closest transmitter. The transmission radii for routing through the net are then chosen to minimize delay under the assumption of relatively light loading. One of the major results is that if η is the arrival rate of new packets to the network, then the optimal transmission radius is proportional to $\eta^{-1/3}$. It is intuitively reasonable that the transmission radius should increase as the loading gets lighter, but the exact way it should increase is rather surprising.

In another S. M. thesis, J. Marcus [Mar84] investigated transmission radii for grid networks and linear array networks. One of the surprising results here was the extreme sensitivity of the results to assumptions about capture effects. In another S.M. Thesis, J. Escobar [Esc85] developed an iterative algorithm for position location of transmitters in a packet radio net.

2.3 Distributed Algorithms

2.3.1 Failure Recovery Algorithms: When links and/or nodes fail in a data network, it is necessary to have a distributed algorithm that notifies other nodes about the failure. If other failures occur while such an algorithm is running, it becomes non-trivial to devise algorithms that always work correctly. Humblet (see [SoH87]) discovered a flaw in

some of the better known algorithms ([Fin79], [Seg83]) and this has reawakened research interest on these algorithms.

J. Spinelli ([Spi85], [Spi86]), in an S.M. thesis, developed a new distributed algorithm for broadcasting topology information in the presence of link and node failures. This algorithm has the interesting property of requiring no sequence numbers or time stamps in the update messages. Because of the simplicity of the algorithm, the proof of correctness is correspondingly simple. A difficulty with this algorithm is in the potentially large number of update messages required when a large number of failures occur. Fortunately, the amount of time required for the algorithm is very small, growing linearly with the sum of the number of nodes and the number of failures and restorals.

I. Castineyra [Cas86], in a Ph. D. thesis, considered the problem of routing in a very unreliable network. The assumption is that links fail and come back up again with such rapidity that algorithms for basing routing decisions on updated topology information become impractical. The investigation thus considered sending packets simultaneously over two routes, attempting to choose the routes to maximize the probability that one or the other would lead to successful reception. The complexity of a number of techniques for finding such pairs of routes were analyzed.

2.3.2 Other Distributed Algorithms and Protocols: One of the classical problems of distributed algorithms for networks is to find the shortest paths from all nodes to a given node in the special case in which all link weights are unity. This is known as the breadth first search problem or the minimum hop problem. The question of interest here is to find an algorithm with the minimum communication complexity (i.e., the worst case number of bits required in update messages) and minimum time complexity (where there might be a tradeoff between communication complexity and time complexity). Awerbuch and Gallager [AwG85a] have developed an algorithm with communication complexity

$O(V^{1.6} + E)$ where V is the number of nodes and E the number of links in the network. For dense networks with $E \geq V^{1.6}$, this order of communication complexity is optimum. The time complexity of the algorithm is $O(V^{1.6})$. The same authors [AwG85b] developed another more complicated algorithm with communication complexity $O(E^{1+\epsilon})$ where ϵ goes to zero very slowly as V approaches infinity.

Another classical problem in the field of asynchronous distributed algorithms is that of electing some given node, say the highest numbered node, in a network as a "leader" given no initial knowledge within the nodes about the identities of the other nodes. Humblet, in [Hum84], gave a particularly simple and efficient solution in the special case of a network with links between each pair of nodes.

In an attempt to understand the relation between noisy communication and distributed algorithms, Gallager [Gal85b] investigated a fully connected broadcast network of $N+1$ nodes in which each node has a single binary digit to communicate. Each binary digit sent by a node is received by all other nodes through independent binary symmetric channels with crossover probability ϵ . Two problems are considered; the first is for a given receiver (the $N+1^{\text{st}}$ node) to reliably find the parity of all the other nodes' information bits, and the other is for the receiver to reliably find all the other nodes' information bits. Distributed algorithms are developed for the nodes to co-operatively solve each of these problems by transmitting $O(\log \log N)$ binary digits. Note that it is easy to solve these problems if each node simply sends its own information bit $O(\log N)$ times, but real co-operation is required to reduce the number of transmissions to $O(\log \log N)$.

REFERENCES

- [Ari85] E. Arikan, "Sequential Decoding for Multiple Access Channels," Ph. D. Thesis, Dept. of EECS, M.I.T., Nov. 1985; also Report LIDS-Th-1517; also published in part in "An Upper Bound to the Cutoff Rate of Sequential Decoding," IEEE Trans. I.T., IT-34, Jan. 1988, pp 65-63.
- [[AwG85a] B. Awerbuch and R. G. Gallager, "Communication Complexity of Distributed Shortest Path Algorithms", M.I.T., Report LIDS-P-1473, June 1985; also IEEE Trans. I.T., vol. IT-33, No. 3, pp. 315-322, May 1987.
- [AwG85b] B. Awerbuch and R.G. Gallager, "Distributed BFS Algorithms," presented at IEEE Conf. Foundations of Computer Science, Portland Ore., 1985.
- [BaE81] D.J. Baker and A. Ephremides, "The Architectural Organization of a Packet Radio Network via a Distributed Algorithm," IEEE Trans. Comm., Vol. COM-29, pp. 1694-1701, Nov. 1981.
- [BeB84] D. P. Bertsekas and D. El Baz, "Distributed Asynchronous Relaxation Methods for Convex Network Flow Problems," LIDS-P-1417.
- [Ber82] D. P. Bertsekas, "Optimal Routing and Flow Control Methods for Communication Networks", in Analysis and Optimization of Systems, A. Bensoussan and J. L. Lions, Eds., New York: Springer-Verlag, 1982, pp. 615-643.
- [Cas86] I. M. Castineyra, "Routing in Unreliable Networks", M.I.T., Department of Electrical Engineering and Computer Science, Ph.D. Thesis, August 1986; also Report LIDS-Th-1577.
- [CMP80] T. M. Cover, R. J. McEliece, and E. C. Posner, "Asynchronous Multi-access channel Capacity," IEEE Trans. I.T., Vol.IT-26, pp. 291-298, Nov. 1980.
- [Esc85] J. Escobar, "Iterative Approach to Position Location in Packet Radio Networks", M.I.T., Department of Electrical Engineering and Computer Science, M.S. Thesis, July 1985; also Report LIDS-TH-1480.
- [Fin79] S.G. Finn, "Resynch Procedures and a Fail-Safe Network Protocol," IEEE Trans. Comm., Vol. COM-27, pp. 840-845, June 1979.
- [Gal77] R. G. Gallager, "A Minimum Delay Routing Algorithm Using Distributed Computation", IEEE Trans. Commun., Vol. COM-25, pp. 73-85, 1977.
- [Gal85a] R. G. Gallager, "A Perspective on Multiaccess Channels," IEEE Trans. I.T., Vol. IT-31, pp. 124-143, Mar. 1985; also M.I.T., Report LIDS-P-1428.
- [Gal85b] R. G. Gallager, "Finding Parity in a Broadcast Network", M.I.T., Report LIDS-P-1490, August 1985; also IEEE Trans. I.T., Mar. 1988, pp176-180.
- [GeP85] L. Georgiadis and P. Papantoni-Kazakos, "A 0.487 throughput limited sensing algorithm," Univ. Conn., Storrs, 1985.

- [HaG86] E. L. Hahne and R. G. Gallager, "Round Robin Scheduling for Fair Flow Control in Data Communication Networks", M.I.T., Report LIDS-P-1537.
- [Hah86] E. L. Hahne, "Round Robin Scheduling for Fair Flow Control in Data Communication Networks", M.I.T., Department of Electrical Engineering and Computer Science, Ph.D. Thesis, December 1986; also M.I.T. Report LIDS-TH-1631.
- [HuH85] J. Y. N. Hui and P. A. Humblet, "The Capacity Region of the Totally Asynchronous Multiple-access Channel," IEEE Trans. I.T., Vol.IT-31, pp. 207-217, Mar. 1985.
- [Hui84a] J.Y.N. Hui, "Multiple Access for the Collision Channel without Feedback," IEEE J. Sel Areas of Comm., vol. SAC-2, No. 4, pp. 575-582, July 1984.
- [Hui84b] J.Y.N. Hui, "Throughput Analysis for Code Division Multiple Accessing of the Spread Spectrum Channel," IEEE J. Sel Areas of Comm., Vol.SAC-2, No. 4, pp. 482-486, July 1984.
- [Hum84] P. A. Humblet, "Selecting a Leader in a Clique in $O(N \log N)$ Messages," LIDS-P1407.
- [Hum85a] P. A. Humblet, "On The Throughput of Channel Access Algorithms with Limited Sensing", M.I.T., Report LIDS-P-1426, January 1985; also IEEE Trans. Comm., vol-COM-34, pp. 345-347, April 1986.
- [Hum85b] P. A. Humblet, "Error Exponents for a Direct Detection Optical Channel," M.I.T. LIDS Report LIDS-P-1337, Oct. 1983; also Proc. 23rd Allerton Conference on Communication, Control and Computing, Monticello, Ill., Oct. 1985.
- [JaB67] I.M. Jacobs and E.R. Berlekamp, "A Lower Bound to the Distribution of Computation for Sequential Decoding," IEEE Trans I.T., April 1967.
- [Kha85] A. C. Khanna, "An Analysis of Multiaccess Reservation Strategies for Satellite Channels", M.I.T., Department of Electrical Engineering and Computer Science, M.S. Thesis, June 1985; also M.I.T. Report LIDS-TH-1471.
- [Mar84] J. N. Marcus, "Throughput for Packet Radio Networks with Dynamic Power Levels", M.I.T., Department of Electrical Engineering and Computer Science, M.S. Thesis, August 1984; also M.I.T. Report LIDS-TH-1390.
- [Mas81] J. L. Massey, "Capacity, Cutoff Rate, and Coding for a Direct Detection Optical Channel," IEEE Trans. Comm., Vol.COM-29, #11, pp 1615-1621, Nov. 1981.
- [McE81] R. J. McEliece, "Practical Codes for Photon Communication," IEEE Trans. I.T., Vol. IT-27, #4, July 1981.
- [Mer84] S. S. Merchant, "The Design and Performance Analysis of an Arbiter for a Multi-Processor Shared-Memory System", M.I.T., Department of Electrical Engineering and Computer Science, M.S. Thesis, September 1984; also M.I.T. Report LIDS-TH-1396.

- [Mos84] J. Mosely, "Asynchronous Distributed Flow Control Algorithms", M.I.T., Department of Electrical Engineering and Computer Science, Ph.D. Thesis, November 1984; also M.I.T. Report LIDS-TH-1415.
- [Muk86] U. Mukherji, "A Schedule-Based Approach for Flow-Control in Data Communication Networks", M.I.T., Department of Electrical Engineering and Computer Science, Ph.D. Thesis, January 1986; also M.I.T. Report LIDS-TH-1527.
- [Par85] A. K. Parekh, "Minimizing the number of Clusters in Mobile Packet Radio Networks", S.M. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., Oct. 1985; also M.I.T. Report, LIDS-TH-1502.
- [Pip81] N. Pippenger, "Bounds on the Performance of Protocols for a Multiple Access Broadcast Channel," IEEE Trans. I.T., Vol. IT-27, No. 2, March 1981
- [Reg86] J. M. Regnier, "Priority Assignment in Integrated Services Networks", M.I.T., Department of Electrical Engineering and Computer Science, Ph.D. Thesis, May 1986; also M.I.T. Report LIDS-TH-1565.
- [ReH87a] J. Regnier and P. Humblet, "Average Waiting Time Assignment Part 1: The Single Link Case", M.I.T. Report LIDS-P-1707, September 1987.
- [ReH87b] J. Regnier and P. Humblet, "Average Waiting Time Assignment Part 2: The Integrated Services Network Case", M.I.T. Report LIDS-P-1708, September 1987.
- [Ros86] C. Rossi, "Transmission Radii for Packet Radio Systems with Capture", S.M. Thesis, Dep't of Electrical Engineering and Computer Science, M.I.T. Sept. 1986.
- [Seg83], A. Segall, "Distributed Network Protocols," IEEE Trans. I.T., Vol. IT-29, No. 1, Jan. 1983.
- [ScS80] M. Schwartz and T. E. Stern, "Routing Techniques Used in Computer Communication Networks", IEEE Trans. Commun., Vol. COM-28, pp. 539-559, 1980.
- [Spi85] J. M. Spinelli, "Broadcasting Topology and Routing Information in Computer Networks", M.I.T. Report LIDS-1470, June 1985.
- [Spi86] J. M. Spinelli, "Broadcasting Topology and Routing Information in Computer Networks", M.I.T. Report LIDS-P-1543, April 1986; revised July 87.
- [TaK84] H. Takagi and L. Kleinrock, "Optimal Transmission Ranges for Randomly Distributed Packet Radio Terminals," IEEE Trans. Comm., Vol. COM-32, No. 3, pp. 246-248, March 1984.
- [Tha84] D. M. Thabit, "Transmission Scheduling to Reduce Convex Delay Costs in Packet Networks," LIDS-TH-1416.
- [Yat86] R. D. Yates, "Methods of Multiple Access Communication with Energy Detectors", M.I.T., Department of Electrical Engineering and Computer Science, S.M. Thesis, June 1986; also M.I.T. Report LIDS-TH-1557.